arXiv:1605.04542v1 [math.ST] 15 May 2016

On stepwise regression

Laurie Davies
Fakulät Mathematik
Universität Duisburg-Essen
laurie.davies@uni-due.de

Abstract

Given data y and k covariates x one problem in linear regression is to decide which in any of the covariates to include when regressing y on the x. If k is small it is possible to evaluate each subset of the x. If however k is large then some other procedure must be use. Stepwise regression and the lasso are two such procedures but they both assume a linear model with error term. A different approach is taken here which does not assume a model. A covariate is included if it is better than random noise. This defines a procedure which is simple both conceptually and algorithmically.

1 Introduction

In a forward stepwise regression the next variable to be included is the one which gives the largest reduction in the sum of squared residuals. The decision as to whether to include this variable is based on the result of an F-test which in turn assumes a linear model

$$Y = x^t\beta + \varepsilon. \tag{1}$$

The f-test does not take the adaptive nature of the procedure into account. Such a test exist for the lasso (Lockhart et al. (2014)) but it is also based on the model (1) and requires assumptions for its validity.

The procedure described below is based on Davies (2016) . It does not assume a model and consequently makes no assumptions about about an error term or about the covariates. It is in other words a procedure. It is legitimate and possible to investigate its behaviour under the model (1) but this will not be done here.

2 The procedure

2.1 Least squares

Suppose that of the k covariates $k_1 < k$ of them have already been included and their sum of squares by $ss(k_1)$. There remain $k_0 = k - k_1$ covariates and the candidate for inclusion is the one whose inclusion decreases the sum of squared residuals by the most. Denote this sum of squared residuals by $ss(k_0)$ so that the reduction in the sum of squares is

$$ss(k_1) - ss(k_0). \tag{2}$$

Replace nor the k_0 covariates by i.i.d. $N(0.1)$ random variables. If one of these is included together with the k_1 covariates already included it is a simple exercise to see that the sum of squared residuals is approximately

$$ss(k_1) - \frac{ss(k_1)}{n}\chi_1^2. \tag{3}$$

Choosing those random variable which lead to the largest reduction in the sum of squares leads to a reduction

$$ss(k_1) - \frac{ss(k_1)}{n}\max\{\chi_1^2, \ldots, \chi_1^2\} = \frac{ss(k_1)}{n}E(k_0) \tag{4}$$

where the maximum is taken over k_0 independent χ_1^2 random variables. The probability that the best of random variables is better than the best of the remaining remaining k_0 covariates is therefore

$$\boldsymbol{P}\left(ss(k1) - \frac{ss(k_1)}{n}E(k_0) < ss(k_1) - ss(k_0)\right) \tag{5}$$

or equivalently

$$\boldsymbol{P}\left(E(k_0) > \frac{n}{ss(k1)}\left(1 - \frac{ss(k_0)}{ss(k_1)}\right)\right). \tag{6}$$

If this probability is reasonably large, say 0.1, then in 10% of the cases the included covariate is no better than random noise. This probability must be specified in advance by a number α. If

$$\boldsymbol{P}\left(E(k_0) > \frac{n}{ss(k1)}\left(1 - \frac{ss(k_0)}{ss(k_1)}\right)\right) < \alpha \tag{7}$$

then the covariate is included. Otherwise the procedure is terminated.

As $E(k_0)$ is the maximum of k_0 χ^2_1 random variables

$$\boldsymbol{P}(E(k_0) > x) = 1 - \text{pchisq}(x, 1)^{k_0} \tag{8}$$

where $\text{pchisq}(x, \nu)$ is the distribution function of a χ^2 random variable with ν degrees of freedom. The covariate is therefore included if

$$\frac{n}{ss(k1)}\left(1 - \frac{ss(k_0)}{ss(k_1)}\right) > \text{qchisq}((1-\alpha)^{1/k_0}, 1). \tag{9}$$

where $\text{qchisq}(x, \nu)$ is the inverse distribution function of a χ^2 random variable with ν degrees of freedom. More informatively one can calculate the P-value

$$1 - \text{pchisq}\left(\frac{n}{ss(k1)}\left(1 - \frac{ss(k_0)}{ss(k_1)}\right)\right)^{k_0} \tag{10}$$

2.2 M-regression

The method can in principle (with the obvious modifications) be applied to L_1 regression but with the disadvantage that there does not exist a simple expression corresponding to (3). If there is a particular interest in L_1 regression then simulations will be required. If however L_1 regression is only used as a protection against outlying y-values this can also be provided by M-regression for which a version of (3) is available.

Let ρ by a symmetric positive twice differentiable convex function with $\rho(0) = 0$. The default function will be the ρ function used in Davies (2014), namely

$$\rho_c(u) = \begin{cases} |u|, & |cu| \geq 15 \\ 2\log(0.5 + 0.5\exp(cu))/c - u, & |cu| < 15 \end{cases} \tag{11}$$

where c is a tuning constant with default value $c = 1$. An alternative choice could be Huber's ρ-function with a tuning constant (Huber and Ronchetti (2009)). The sum of squared residuals $ss(k)$ is replaced by

$$s_\rho(k) = \text{argmin}_\beta \sum_{i=1}^{k} \rho(y_i - x_i^t\beta). \tag{12}$$

As it stands $s_\rho(k)$ is not satisfactory and must be augmented by a data dependent scale value σ to give

$$s_\rho(k, \sigma) = \text{argmin}_\beta \sum_{i=1}^{k} \rho\left(\frac{y_i - x_i^t\beta}{\sigma}\right). \tag{13}$$

For a given ρ and σ $s_\rho(k,\sigma)$ can be calculated using the algorithm described in **7.8.2** of Huber and Ronchetti (2009). Typically only a few number of iterations are required.

Given all this (3) is replaced by

$$s_\rho(k_1,\sigma) - \frac{\sum_{i=1}^n \rho^{(1)}\left(\frac{r_i}{\sigma}\right)^2}{\sum_{i=1}^n \rho^{(2)}\left(\frac{r_i}{\sigma}\right)}\chi_1^2 = s_\rho(k_1,\sigma) - \frac{s_{\rho^{(1)}}(k_1,\sigma)}{s_{\rho^{(2)}}(k_1,\sigma)}\chi_1^2 \tag{14}$$

where $\rho^{(1)}$ and $\rho^{(2)}$ are th first and second derivatives of ρ respectively,

$$r_i = y_i - x_i^t\hat{\beta} \quad \text{with} \quad \hat{\beta} = \text{argmin}_\beta \sum_{i=1}^{k_1} \rho\left(\frac{y_i - x_i^t\beta}{\sigma}\right)$$

and

$$s_{\rho^{(1)}}(k_1,\sigma) = \sum_{i=1}^n \rho^{(1)}\left(\frac{r_i}{\sigma}\right)^2, \quad s_{\rho^{(2)}}(k_1,\sigma) = \sum_{i=1}^n \rho^{(2)}\left(\frac{r_i}{\sigma}\right).$$

The P-value (10) becomes

$$1 - \text{pchisq}\left(\frac{s_{\rho^{(1)}}(k_1,\sigma)}{s_{\rho^{(2)}}(k_1,\sigma)}\left(1 - \frac{s_\rho(k_0,\sigma)}{s_\rho(k_1,\sigma)}\right)\right)^{k_0} \tag{15}$$

It remains to specify the choice of scale σ. The procedure described here uses a σ dependent on the k_1 variables already incorporated. This same σ is used to judge whether a new variable is to be included. This is why there is only one value of σ in (15). One possibility is to do a full M-regression and for both location and scale based on the k_1 covariates and take σ to be the scale part (Huber and Ronchetti (2009)). This has a certain intellectual consistency but at the expense of greater programming effort. Instead the following procedure will be used. If a new covariate is to be included then the residuals r_i are calculated from an M-regression using the $k_1 + 1$ covariates but based on the σ for the original k_1 covariates. The new σ is taken to be the median absolute deviation of the r_i multiplied by the Fisher consistency factor 1.48 which is the default version of the MAD in `R`. The procedure is started using the residuals from best L_1 single covariate calculated using for example Koenker (2010).

3 Two examples

The method will be illustrated using the prostate cancer data also used in (Lockhart et al. (2014)) and the low birth weight data from Hosmer and Lemeshow (1989).

	P-value	
covariate	L_2	M
lcavol	0.0000	0.0000
lweight	0.0122	0.0083
svi	0.0123	0.0101
lbph	0.4233	0.3408
age	0.4952	0.4083
pgg45	0.5541	0.4839
lcp	0.4093	0.2845
gleason	0.7636	0.7300

Table 1: The prostate data: the covariates in order of inclusion and their P-values

The prostate cancer data were obtained from Lokhorst et al. (2014). They are described in Hastie et al. (2008). The sample size is $n = 97$ with eight covariates. Table 1 gives the order in which the covariates entered the regression together with their P-values for the L_2 and M regressions. The order was the same for both. Table 2 is the same but with the first y value changed from -0.4307829 to 10. It shows that the results for the M regression remain stable but those for the L_2 regression change considerably apart from the covariate *lcavol*.

The dependent variable in the low birth weight data is taken to be the weight of the child. The covariates are:
(1) Age of mother, (2) Weight of mother, (3) Smoking status, (4) History of premature labor, (5) History of hypertension, (6) Uterine irritability, (7) Number of physician visits, (8) Race-1, (9) Race-2.
Model and functional choice for this data set has been considered in Davies (2014) and Claeskens and Hjort (2003) (model choice) and Davies (2016) (functional choice). Table 3 gives the results for the stepwise functional choice. The oder of the covariates is the same for both methods. The choice (6,9,3) with $\alpha = 0.05$ corresponds to the functional encoded as 292 in Davies (2016) which one of the functionals chosen after considering all subsets with $\alpha = 0.05$. The

subset (6,9,3,5) corresponds to the functional encoded as 308 with $\alpha = 0.1$ in Davies (2016).

References

Claeskens, G. and Hjort, N. L. (2003), 'Focused information criterion', *Journal of the American Statistical Association* **98**, 900–916.

Davies, L. (2014), *Data Analysis and Approximate Models*, Monographs on Statistics and Applied Probability 133, CRC Press.

Davies, P. L. (2016), 'Functional choice and non-significance regions in regression', arXiv:1605.01936 [math.ST].

Hastie, T., Tibshirani, R. and Friedman, J. (2008), *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, second edn, Springer, New York.

Hosmer, D. W. and Lemeshow, S. (1989), *Applied Logistic Regression*, Wiley, New York.

Huber, P. J. and Ronchetti, E. M. (2009), *Robust Statistics*, second edn, Wiley, New Jersey.

Koenker, R. (2010), 'quantreg: Quantile regression', http://CRAN.R-project.org/package=quantreg. R package version 4.53.

Lockhart, R., Taylor, J., Tibshirani, R. J. and Tibshirani, R. (2014), 'A significance test for the lasso', *Ann. Statist.* **42**(2), 413–468.

Lokhorst, J., Venables, W. N. and Turlach, B. (2014), 'lasso2', http://www.maths.uwa.edu.au/∼berwin/software/lasso.html. R-package.

L_2		M	
covariate	P-value	covariate	P-value
lcavol	0.0000	lcavol	0.0000
svi	0.1234	svi	0.0176
age	0.6623	lweight	0.0366
lbph	0.4534	lbph	0.4676
lweight	0.950	age	0.1766
pgg45	0.7615	pgg45	0.5309
lcp	0.7615	lcp	0.3337
gleason	0.8949	gleason	0.8269

Table 2: The prostate data but with $y(1) = 10$: the covariates in order of inclusion and their P-values

	P-value	
covariate	L_2	M
6	0.0009	0.0008
9	0.0187	0.0223
3	0.0015	0.0009
5	0.0934	0.1017
2	0.0778	0.0649
8	0.8842	0.8616
4	0.9285	0.9038
1	0.8779	0.8359
7	0.7557	0.7607

Table 3: The low birth weight data: the covariates in order of inclusion and their P-values.